


The 2% Solution

Anton Zhukov



Co-organizer

Yandex

Who am I?

- Infrastructure Engineer
- Platform Team TechLead
- 12+ years in development
- 4 years in ManyChat
- Specialization: high-loaded systems, performance engineering and fail-safe solutions



Manychat

What is ManyChat

1M+

Businesses
chose
ManyChat
for growth

190+

Countries use
ManyChat
across the
world

1B+

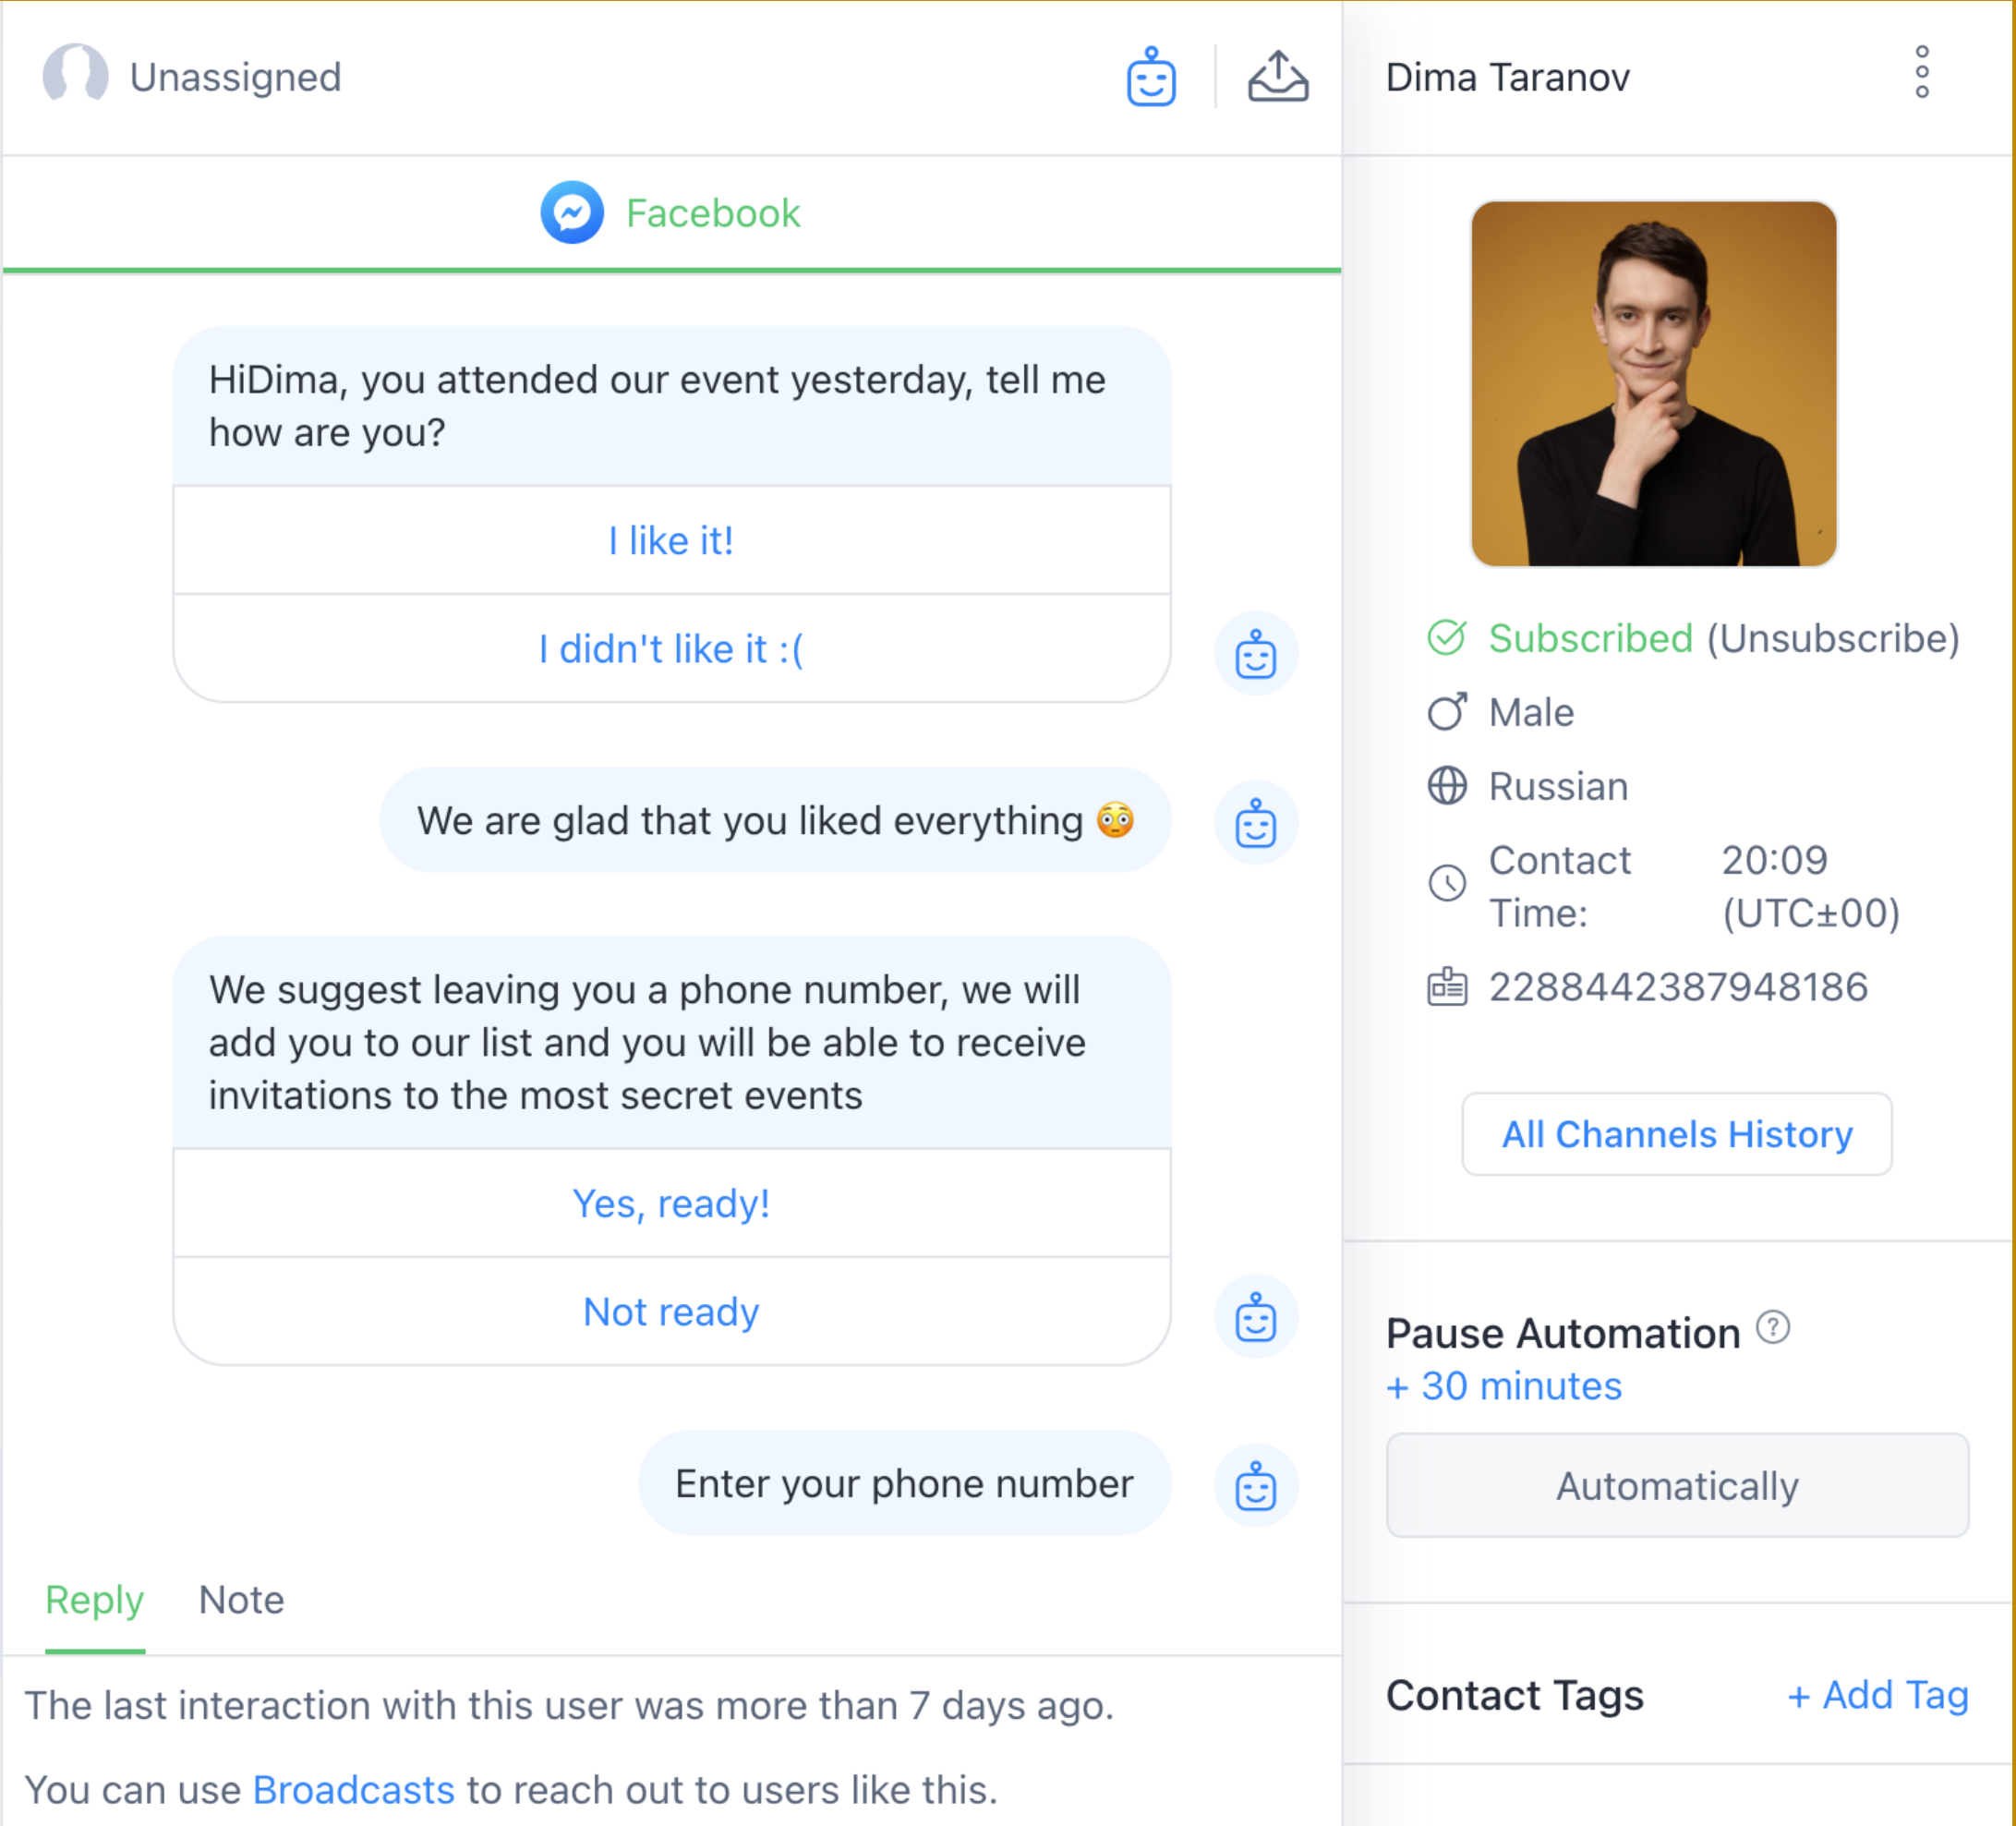
Conversations
powered
by ManyChat
every year

140+

People are
in ManyChat
team

Usecases of events DB

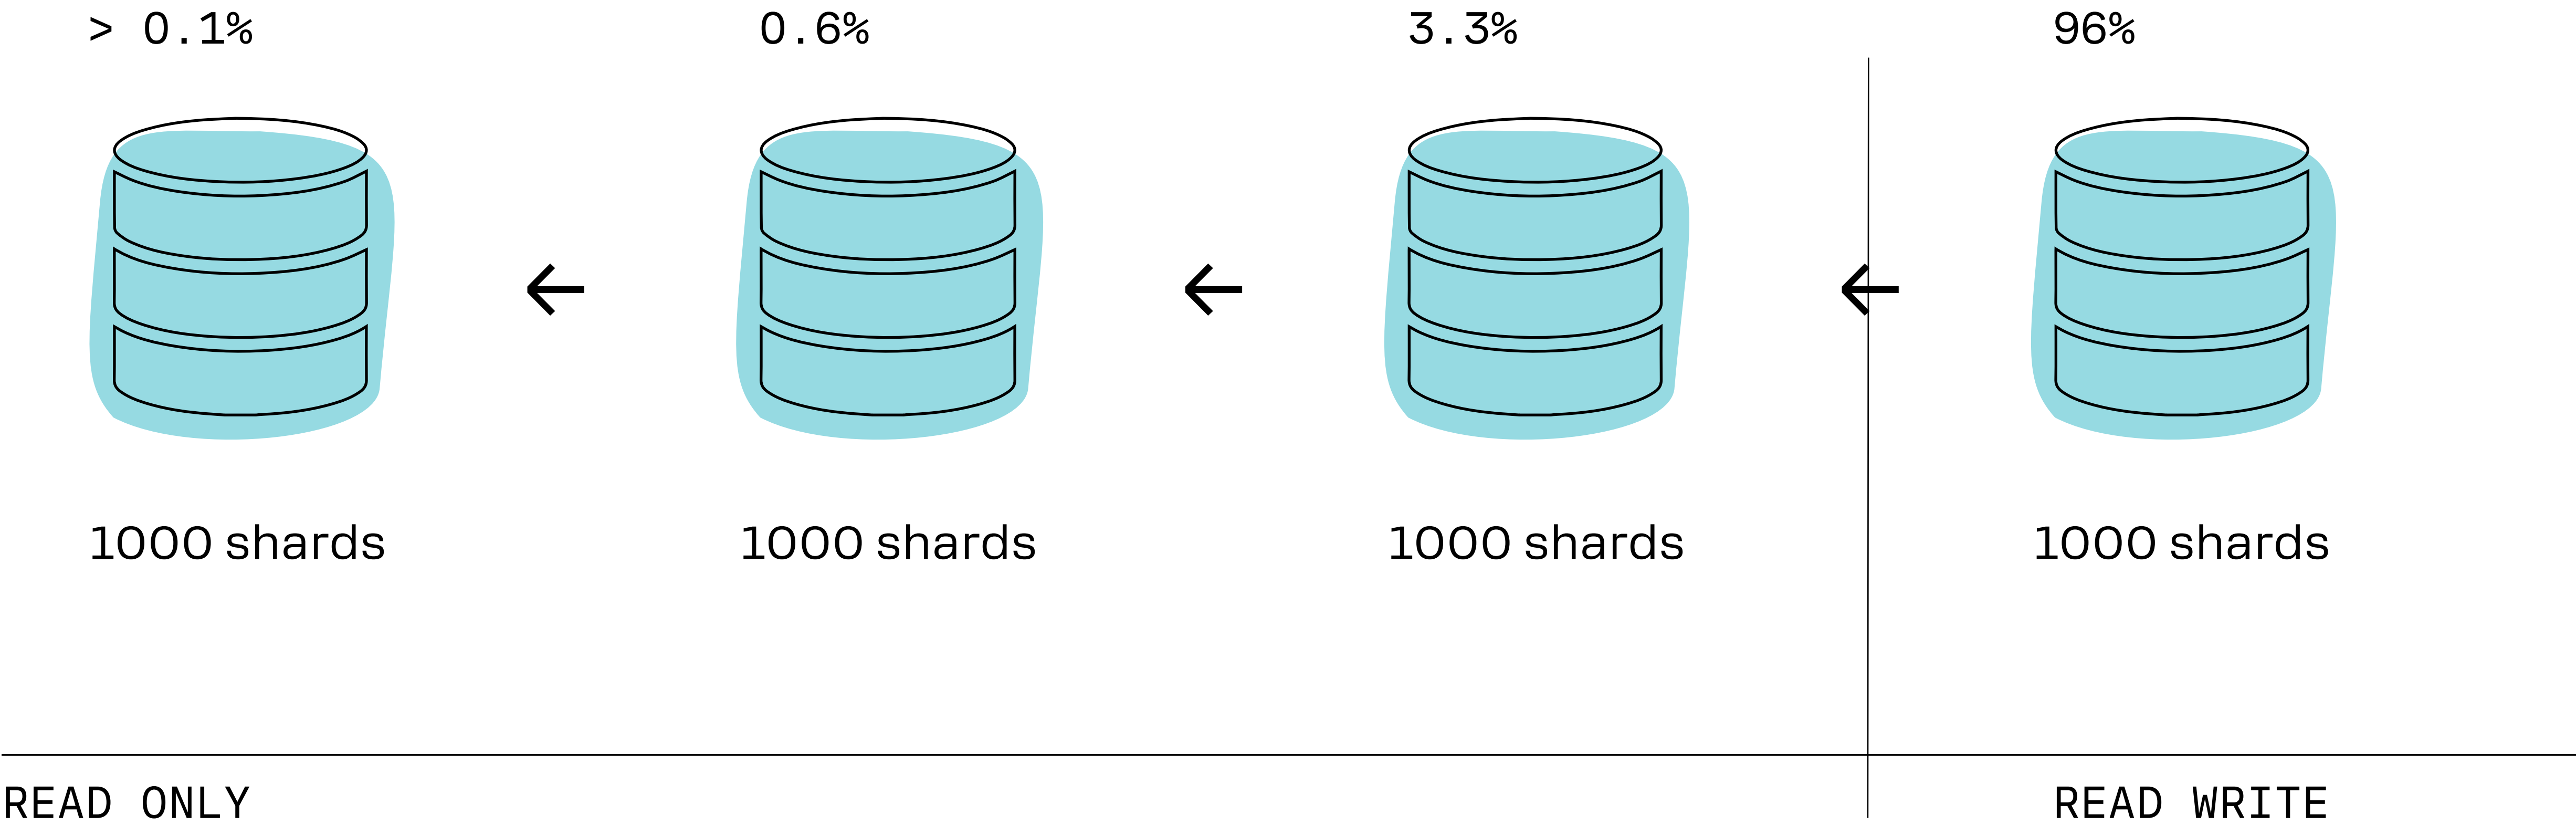
- 1 Messages
- 2 Actions (external / internal)
- 3 Historical state storage



The Trouble

- 1 PostgreSQL Database with events
- 2 64TB of data
- 3 Events are cold and rarely read
- 4 10 000\$ monthly price...

Our event databases lineup



YES

**REMOVE
USELESS DATA?**

BUT

**LOW
IMPACT**

YES

REMOVE OLD
DATA OF FREE
ACCOUNTS?

Discovering a solution

BUT

A LITTLE
DIRTY AND
QUESTIONABLE
WAY

YES

**CHANGE
THE DATABASE?**


BUT

**DO WE HAVE
ENOUGH TIME TO
SHARE EXPERTISE
AND COLLECT
BUGS?**


Discovering a solution

But, what do we want actually?

DECREASE COSTS!



DECREASE THE DATA SIZE




CHANGE THE DATA STORAGE

Our targets

- 1 Parallel stateless component
- 2 Fast access by scope (index similarity)
- 3 Fast and stable
- 4 Encryption
- 5 Compression
- 6 Ability to rollback (this is an experiment)
- 7 Low costs of storage
- 8 Low costs of support
- 9 Repeatable case (that's not only a single database)
- 10 Very low price

Change event databases lineup



1. Last in a call chain
(not required)

2. Empty resultset while a
donor is in migration

3. When a donor is disabled, the
cold storage begins to work
automatically

Discovering a solution

1 Major slice of event
count per subscriber

2 Chunk
size

3 Upload /
Download costs


4 Compression /
Decompression costs

Metrics for making a decision

5 Encryption /
Decryption costs

6 Total price including
development, storage
and usage

Main slice migration



Discovering a solution

The choice of metrics:

- 1 ≥ 70 events
to cold storage (70 - 2500)
- 2 < 70 events for direct
migration
- 3 Brotli for
compression
- 4 AES-256-CBC for encryption
- 5 AWS S3 as a storage

Main Aspects of the Migration Process

- 1 Have migration data
- 2 Step by step migration (chunk → compression → encryption → upload to s3 → insert link events to a database)
- 3 Database should be really “cold”
- 4 Remember about database capacity
- 5 Consistency control

Main Aspects of the Migration Process

account	subscriber	events	status	migration	step_data	metrics
1	1	70	created	1	{“some_data”: 1}	{“some_metrics”: 1}
1	2	12	chunk_created	2	{“some_data”: 1}	{“some_metrics”: 1}
1	3	2320	encrypted	1	{“some_data”: 1}	{“some_metrics”: 1}
1	4	4	cleaned	2	{“some_data”: 1}	{“some_metrics”: 1}
2	5	70	compressed		{“some_data”: 1}	{“some_metrics”: 1}
2	6	70	events_created		{“some_data”: 1}	{“some_metrics”: 1}

Main Aspects of the Migration Process

Step by step

[2022-11-28 18:56:16.00]: db: 62% compress: 3% encrypt: 4% upload: 27% links: 4% clean: 0.2% total time: 12.2s

[2022-11-28 18:56:16.00]: db: 57% compress: 3% encrypt: 4% upload: 32% links: 4% clean: 0.1% total time: 34.1s

[2022-11-28 18:56:16.00]: db: 79% compress: 3% encrypt: 4% upload: 10% links: 4% clean: 0.2% total time: 13.1s


[2022-11-28 18:56:16.00]: db: 2% compress: 3% encrypt: 4% upload: 87% links: 4% clean: 0.2% total time: 17.6s

[2022-11-28 18:56:16.00]: db: 12% compress: 3% encrypt: 54% upload: 27% links: 4% clean: 0.2% total time: 54.9s

[2022-11-28 18:56:16.00]: db: 12% compress: 53% encrypt: 4% upload: 27% links: 4% clean: 0.2% total time: 51.2s

Main Aspects of the Migration Process

Selecting rows



Main Aspects of the Migration Process

Consistency control

- 1 You can check everything... however with a longer process than migration
- 2 Or, check it by numbers
- 3 Live checking based on a sample

Main Aspects of the Migration Process

Switch shards and disable a database

- 1 The easiest... and the most scary
- 2 Switch shards by flag
- 3 Shutdown a database
- 4 Drop a database

How does it work

Link events

DIRECT

```
{  
  "s3Key":  
    "event/4000/97/300897/4000_300897_9427.data.  
    br.enc",  
  "first_event_id": 3823,  
  "min_event_timestamp": "2019-11-29 03:10:43.6",  
  "max_event_timestamp": "2020-05-23 14:35:37.8",  
  "size": 136  
}
```

EVENT_ID = 9427


REVERSE

```
{  
  "s3Key":  
    "event/4000/97/300897/4000_300897_9427.data.  
    br.enc",  
  "last_event_id": 9427,  
  "min_event_timestamp": "2019-11-29 03:10:43.6",  
  "max_event_timestamp": "2020-05-23 14:35:37.8",  
  "size": 136  
}
```


EVENT_ID = 3823

How does it work

Engine




→ → →



DESC

How does it work


Engine



ASC

Post factum decisions

Hot accounts



Post factum decisions

- 1 Event source
- 2 Runtime consistency control

Design mistakes

- 1 S3 chunk address without db id
- 2 Migration to a production database
- 3 And the one little mistake

The one little mistake

- 1 Difference between tests and reality
- 2 Always look at metrics
- 3 Re-migration price
- 4 But... we have the one big plus...

What do we have now

1 Simple, stable and cheap storage

2 Decreased annual costs from 427 200\$ to 8 544\$

3 Total saved 500 000\$+

4 Scalable and repeatable case

5 Equal uptime comparing only with only PostgreSQL database

6 Freedom from useless expertise sharing and unpredictable point of failure

7 Keeping fidelity to the principle “Do it simple”

Metrics

	BEFORE	1 DB IN CS	3 DB'S IN CS
DATA SIZE	226TB	9.75TB	34TB
FIRST REQUEST TIME	60-82MS	60-82MS	60-82MS
STORAGE PRICE PER MONTH	35 600\$	190\$	712\$

YES

COMPRESS MORE

What could be done but doesn't make sense

BUT

190\$

The major point — costs of development and support, we decreased costs

YES

MERGE SMALL
CHUNKS
OF SEVERAL
SUBSCRIBERS

What could be done but doesn't make sense

BUT

190\$

The major point — costs of development
and support, we decreased costs

YES

NORMALIZE
EQUAL EVENTS
AND DECREASE
EVENT COUNT

What could be done but doesn't make sense

BUT

190\$

The major point — costs of development and support, we decreased costs

Thank you!

TELEGRAM:
@ANTON_D_ZHUKOV


TWITTER:
@ANTONDZHUKOV

MAIL:
AJ@MANYCHAT.COM



Leave your feedback!

You can rate the talk and give
a feedback on what you've
liked or what could be
improved



Co-organizer

Yandex